

A Near-Linear Pseudopolynomial Time Algorithm for Subset Sum

Karl Bringmann*

October 18, 2016

Abstract

Given a set Z of n positive integers and a target value t , the SUBSETSUM problem asks whether any subset of Z sums to t . A textbook pseudopolynomial time algorithm by Bellman from 1957 solves SUBSETSUM in time $\mathcal{O}(nt)$. This has been improved to $\mathcal{O}(n \max Z)$ by Pisinger [J. Algorithms'99] and recently to $\tilde{\mathcal{O}}(\sqrt{n}t)$ by Koiliaris and Xu [SODA'17].

Here we present a simple and elegant randomized algorithm running in time $\tilde{\mathcal{O}}(n+t)$. This improves upon a classic algorithm and is likely to be near-optimal, since it matches conditional lower bounds from SETCOVER and k -CLIQUE.

We then use our new algorithm and additional tricks to improve the best known polynomial space solution from time $\tilde{\mathcal{O}}(n^3t)$ and space $\tilde{\mathcal{O}}(n^2)$ to time $\tilde{\mathcal{O}}(nt)$ and space $\tilde{\mathcal{O}}(n \log t)$, assuming the Extended Riemann Hypothesis. Unconditionally, we obtain time $\tilde{\mathcal{O}}(nt^{1+\varepsilon})$ and space $\tilde{\mathcal{O}}(nt^\varepsilon)$ for any constant $\varepsilon > 0$.

*Max Planck Institute for Informatics, Saarland Informatics Campus, Germany, kbringma@mpi-inf.mpg.de.

1 Introduction

Given a (multi-)set Z of n positive integers and a target t , the SUBSETSUM problem asks whether there is a subset Y of Z summing to exactly t . This classic NP-hard problem draws its significance from the fact that it lies at the core of many other NP-hard (optimization) problems, e.g. SUBSETSUM can easily be reduced to the KNAPSACK problem, CONSTRAINED SHORTEST PATH (also known as RESTRICTED SHORTEST PATH), and many other problems.

Bellman showed in 1957 that SUBSETSUM can be solved in pseudopolynomial time $\mathcal{O}(nt)$ by dynamic programming [5]. This algorithm is being taught for decades in undergraduate algorithms courses and thus had a great influence on computer science. Since this pseudopolynomial time algorithm is a fundamental part of our curriculum, and SUBSETSUM is one of the core NP-hard problems, it is an important question whether the running time of $\mathcal{O}(nt)$ can be improved.

Pisinger [22] used basic RAM parallelism (allowing to operate on $\Theta(\log t)$ bits in constant time) to obtain the first improvement: an $\mathcal{O}(nt/\log t)$ algorithm. Considering $s := \max Z$, Pisinger [21] presented an $\mathcal{O}(ns)$ algorithm, which is faster than $\mathcal{O}(nt)$ in some situations. The first algorithm breaking through the $\mathcal{O}(nt)$ barrier by a polynomial factor in the worst case was the recent $\tilde{\mathcal{O}}(\sqrt{n}t)$ algorithm¹ by Koiliaris and Xu [17]. Roughly speaking, they use hashing to solve SUBSETSUM quickly if Z is contained in a small interval. Then they reduce the general case to the small interval case by appropriately splitting Z into smaller sets. They also presented an $\tilde{\mathcal{O}}(t^{4/3})$ algorithm.

There is reason to believe that SUBSETSUM has no $t^{1-\varepsilon}n^{\mathcal{O}(1)}$ algorithm for any $\varepsilon > 0$, since this would yield an $2^{(1-\varepsilon')n}(nm)^{\mathcal{O}(1)}$ algorithm for SETCOVER (on m sets over a universe of size n) via the reductions in [7], and thus break the SETCOVER hypothesis. A second reason is that any combinatorial $t^{1-\varepsilon}n^{\mathcal{O}(1)}$ algorithm for SUBSETSUM would yield a combinatorial $\mathcal{O}(n^{(1-\varepsilon')k})$ algorithm for k -CLIQUE, for some large constant k , via the reduction in [1]². This means that polynomial improvements over running time t , even in terms of n , are unlikely. However, these arguments give no evidence that the additional factor n of the $\mathcal{O}(nt)$ dynamic programming algorithm is necessary (or the factor \sqrt{n} of Koiliaris and Xu). Specifically, they leave the open problem whether there is an $\tilde{\mathcal{O}}(n+t)$ algorithm.

Note that an $\tilde{\mathcal{O}}(n+t)$ algorithm would also up to polylogarithmic factors dominate the $\mathcal{O}(ns)$ algorithm by Pisinger [21], where $s = \max Z$, since any non-trivial instance satisfies $s \geq t/n$.

Near-linear Time We present an algorithm in time $\tilde{\mathcal{O}}(n+t)$. This improves the classic dynamic programming solution by a factor $\tilde{\Omega}(n)$ and the best known algorithm by a factor $\tilde{\Omega}(\sqrt{n})$. Moreover, we match the conditional lower bounds, so any further polynomial improvement would beat the $\mathcal{O}^*(2^n)$ SETCOVER algorithm and possibly the $\mathcal{O}(n^k)$ combinatorial k -CLIQUE algorithm.

Theorem 1.1 (Section 3). *SUBSETSUM can be solved in time $\tilde{\mathcal{O}}(n+t)$ by a randomized, one-sided error algorithm with error probability $(n+t)^{-\Omega(1)}$.*

More precisely, if $\mathcal{S}(Z;t)$ is the set of all sums generated by subsets of Z and bounded by t , then we can compute in time $\mathcal{O}(n+t \log t \log^3(n/\delta) \log n)$ a set $S \subseteq \mathcal{S}(Z;t)$ containing any $s \in \mathcal{S}(Z;t)$

¹For a running time T depending on the input size n and possibly more parameters, such as t , we write $\tilde{\mathcal{O}}(T)$ as shorthand for $\mathcal{O}(T \text{polylog}(T))$. Similarly, $\tilde{\Omega}(T)$ denotes $\Omega(T/\text{polylog}(T))$.

²Abboud et al. [1] presented a reduction from k -CLIQUE to $\mathcal{O}(k^2)$ -SUM on $\mathcal{O}(f(k)n^2)$ integers in the range $\{1, \dots, \mathcal{O}(f(k)(n^{1+o(1)})^k)\}$ for some (explicit) function f . Note that k' -SUM can be reduced to SUBSETSUM by introducing $\mathcal{O}(\log k')$ leading bits that ensure choosing exactly k' integers. This yields a reduction from k -CLIQUE to SUBSETSUM on $\mathcal{O}(f'(k)n^2)$ integers in $\{1, \dots, \mathcal{O}(f'(k)(n^{1+o(1)})^k)\}$ for some function f' . Let $\varepsilon > 0$, $c \in \mathbb{R}$ be constants and let k be sufficiently large ($k \geq 8c/\varepsilon$). Then for all sufficiently large $n \geq n_0$, such that the $n^{o(1)}$ -factor is less than $n^{\varepsilon/4}$, a combinatorial $\mathcal{O}(t^{1-\varepsilon}n^c)$ algorithm for SUBSETSUM would yield a combinatorial $\mathcal{O}(n^{k-\varepsilon k/2})$ algorithm for k -CLIQUE, refuting the combinatorial k -CLIQUE conjecture.

with probability at least $1 - \delta$.

We briefly also consider the UNBOUNDEDSUBSETSUM problem, where each input number may be used multiple times instead of just once, and present a simple deterministic $\tilde{O}(n + t)$ algorithm. In Section 5 we prove this result and discuss why it is much simpler than our main result.

Theorem 1.2. UNBOUNDEDSUBSETSUM can be solved deterministically in time $\mathcal{O}(n + t \log t)$.

Polynomial Space A surprising, relatively recent result is that SUBSETSUM can be solved in pseudopolynomial time and *polynomial* space: Lokshtanov and Nederlof [19] presented an algorithm in time $\tilde{O}(n^3 t)$ and space $\tilde{O}(n^2)$. In contrast, the $\mathcal{O}(nt)$ dynamic programming algorithm as well as all improvements need space $\tilde{\Theta}(t)$, which can be much larger than $\text{poly}(n)$. We use our new SUBSETSUM algorithm and some additional tricks to improve upon their bounds. This almost answers an open problem by J. Nederlof [14] for a time $\tilde{O}(nt)$ and space $\tilde{O}(n)$ algorithm.

Theorem 1.3 (Section 4). SUBSETSUM has a randomized, one-sided error algorithm with error probability $(n + t)^{-\Omega(1)}$ in

- time $\tilde{O}(nt)$ and space $\tilde{O}(n \log t)$ assuming the Extended Riemann Hypothesis (ERH), and
- time $\tilde{O}(n t^{1+\varepsilon})$ and space $\tilde{O}(n t^\varepsilon)$ for any constant $\varepsilon > 0$ unconditionally.

We leave it as an open problem to find an algorithm with time $\tilde{O}(n + t)$ and space $\tilde{O}(n \log t)$.

1.1 Techniques

Our techniques are very different from previous improvements for SUBSETSUM. In particular, all previous algorithms are deterministic. Our near-linear time algorithm is simple and elegant and makes extensive use of randomization and the fast Fourier transform (FFT). In the following we discuss the main ingredients of our algorithms.

Sumset Computation For sets of integers A, B we define the *sumset* $A \oplus B$ as the set of all sums $a + b$ with $a \in A \cup \{0\}$, $b \in B \cup \{0\}$. Note that here we allow to not choose any value in one or both sets by adding $\{0\}$. Often we consider the *t-capped sumset* $A \oplus_t B$, which is simply the restriction of $A \oplus B$ to $\{0, \dots, t\}$. Sumset computation is the most important primitive of our algorithm. Easy reductions transform this problem to Boolean convolution and further to integer multiplication, which has well-known algorithms using FFT. This yields an algorithm for computing $A \oplus_t B$ in time $\mathcal{O}(t \log t)$ (see Section 2 for details).

Color-Coding Color-coding is an algorithmic technique that was first used for the k -PATH problem: Given a graph G , decide whether it contains a path of length k [3]. The idea is to randomly color the vertices of G with k colors, so that for a fixed path of length k in G with probability $1/k!$ it is colored $(1, 2, \dots, k)$, in which case we can find it by a simple dynamic programming algorithm on the layered graph obtained from keeping only the edges in G from color class i to $i + 1$ (for each i). Over $\mathcal{O}(k! \log n)$ repetitions we find a k -path with high probability, if one exists. Research on color-coding has led to various improvements and derandomizations of this technique [3, 20, 23]. The derandomizations also apply to our use of color-coding in this paper, however, other parts of our algorithm seem impossible to derandomize and we leave this as an open problem.

We make use of color-coding for finding all sums generated by small subsets. More precisely, given a SUBSETSUM instance (Z, t) and a threshold k , we compute a set $S \subseteq \mathcal{S}(Z; t)$ containing any sum generated by a subset $Y \subseteq Z$ of size $|Y| \leq k$ with constant probability. This can be boosted

to any high probability by repeating and taking the union. The main trick is to randomly partition $Z = Z_1 \cup \dots \cup Z_{k^2}$ by assigning to any element $z \in Z$ a color in $\{1, \dots, k^2\}$ independently and uniformly at random. Consider the sumset $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$. Note that this set consists of sums generated by subsets of Z , in particular we never add up the same number twice since the Z_i are disjoint. We say that the random partition *splits* Y if each Z_i contains at most³ one element of Y . In this case, the sum $\Sigma(Y) := \sum_{y \in Y} y$ is contained in the sumset $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$. Indeed, by the definition of \oplus_t , in each position i we can choose a number in $Z_i \cup \{0\}$, so if Z_i contains exactly one element of Y then we can choose this element, while if Z_i contains no element of Y we can choose 0, to write $\Sigma(Y)$ as a sum in the sumset $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$. It remains to argue that the random partition splits Y with constant probability. The color-coding assigns each element in Y a random color in $\{1, \dots, k^2\}$, and thus we can view the partitioning, restricted to Y , as placing k balls (the elements of Y) into k^2 bins (the subsets Z_i). This process is well understood, in particular the birthday paradox implies for our choice of k^2 bins that with constant probability some Z_i contains more than one element of Y – which is a bad event in our situation. However, since the bound of the birthday paradox is tight, also with constant probability we have $|Z_i \cap Y| \leq 1$ for all i , and thus the partitioning splits Y . Note that the running time for randomly partitioning and computing the sumset $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$ is $\tilde{O}(n + k^2 t)$, which is near-linear in $n + t$ if k is polylogarithmic in n and t .

Layer Splitting Any SUBSETSUM instance (Z, t) can be partitioned into $\log n$ layers $Z_i \subseteq [t/2^i, t/2^{i-1}]$, plus a set $Z_0 \subseteq [0, t/n]$ that can be treated very similarly to layers and that we ignore here for simplicity. We obtain the desired sumset $\mathcal{S}(Z; t)$ by combining the sumsets $\mathcal{S}(Z_i; t)$ of the layers in a straight-forward way using sumset computations.

It remains to compute $\mathcal{S}(Z_i; t)$. By the definition of layers, any set $Y \subseteq Z_i$ summing to at most t has size $|Y| \leq 2^i$. Thus, the color-coding algorithm with $k = 2^i$ computes $\mathcal{S}(Z_i; t)$, and it runs in time $\tilde{O}(n + tk^2)$. To obtain an $\tilde{O}(n + t)$ algorithm we need to eliminate the factor k^2 . We remove one factor k by observing that all items in Z_i are bounded by $\mathcal{O}(t/k)$, which allows us to implement the sumset computations in the color-coding algorithm more efficiently. The remaining factor k stems from color-coding partitioning Z_i into k^2 subsets. We remove this factor by a two-stage approach: In the first stage, we partition Z_i into roughly $k = 2^i$ sets $Z_{i,j}$. This k is too small to split Y entirely, i.e., to have $|Z_{i,j} \cap Y| \leq 1$ for all j with high probability. Indeed, for this property we would need to partition Z_i into k^2 sets. However, we still have $|Z_{i,j} \cap Y| \leq \mathcal{O}(\log n)$ with high probability. Hence, in the second stage we can run the color-coding algorithm with size bound $k' = \mathcal{O}(\log n)$ on each $Z_{i,j}$, and then combine their computed sumsets in a straight-forward way. This removes the factor- k overhead from partitioning into k^2 sets. Carefully implementing these ideas yields time $\tilde{O}(n + t)$.

Polynomial Space The polynomial space algorithm by Lokshtanov and Nederlof interprets a SUBSETSUM algorithm as a circuit, where each gate performs the convolution or pointwise addition of two vectors of some fixed length $f(n, t)$. This circuit C is transferred to the Fourier domain by replacing every convolution gate by pointwise multiplication. The new circuit C' computes the Fourier transform of the output vector of C . Using that the inverse Fourier transform can be written as a simple sum, we can evaluate an entry of C by evaluating all entries of C' one-by-one. Writing $g(n, t)$ for the number of gates in C , since all operations in C' are pointwise, computing an entry of C' can be done in $\tilde{O}(g(n, t))$ arithmetic operations, storing $\tilde{O}(g(n, t))$ numbers. Evaluating all entries of C' one-by-one thus can be done in $\tilde{O}(g(n, t) \cdot f(n, t))$ arithmetic operations, storing $\tilde{O}(g(n, t))$ numbers.

³In contrast to typical uses of color-coding, we want *at most* one element instead of *exactly* one element.

One problem in the algorithm by Lokshtanov and Nederlof is that entries can become as large as $2^{\Omega(n)}$, so they need to work with $\tilde{\mathcal{O}}(n)$ bits of precision. This yields total time $\tilde{\mathcal{O}}(n \cdot g(n, t) f(n, t))$ and space $\tilde{\mathcal{O}}(n \cdot g(n, t))$. We work with a variant of their algorithm using modular arithmetic instead of complex numbers. This allows us to work *modulo a random prime* p , thus reducing the precision from $\tilde{\mathcal{O}}(n)$ to $\mathcal{O}(\log p)$ bits, and improving time and space by a factor $\tilde{\mathcal{O}}(n/\log p)$. A technical difficulty is that for using the Fourier transform the field \mathbb{Z}_p has to contain a primitive t -th root of unity. This is guaranteed by choosing p in the arithmetic progression $1 + t \cdot \mathbb{N}$. However, to be able to choose p at random from a sufficiently large ground set, we need to choose a threshold x such that there are many primes $p \leq x$ in the arithmetic progression $1 + t \cdot \mathbb{N}$. This requires a quantitative version of Dirichlet's theorem. The best such result assumes ERH and allows us to choose $p \leq x = n^{\mathcal{O}(1)}$, yielding an improvement factor $\tilde{\mathcal{O}}(n/\log p) = \tilde{\mathcal{O}}(n)$.

The specific SUBSETSUM circuit of Lokshtanov and Nederlof needs $g(n, t) = \mathcal{O}(n)$ gates and vector length $f(n, t) = \tilde{\mathcal{O}}(nt)$. A priori it is not easy to improve $f(n, t)$ to $\tilde{\mathcal{O}}(t)$. We show that the circuit induced by our new SUBSETSUM algorithm allows us to set $f(n, t) = \tilde{\mathcal{O}}(t)$ and $g(n, t) = \tilde{\mathcal{O}}(n)$, thus improving the running time by another factor $\tilde{\mathcal{O}}(n)$.

1.2 Further Related Work

SUBSETSUM has been studied extensively, see e.g. [16]. The best-known running time in terms of n is $\mathcal{O}^*(2^{n/2})$ [13]. There is a large literature improving this running time for inputs that are in some sense “structured”, see, e.g., [4] and the references therein. Using number-theoretic arguments, certain *dense* cases of SUBSETSUM are solvable in near-linear time, e.g., if $t \ll n^2$ then there is an $\tilde{\mathcal{O}}(n)$ algorithm [11]. A $(1 - \varepsilon)$ -approximation for SUBSETSUM yields a set with sum in $[(1 - \varepsilon)t, t]$ if one exists. The best known approximation algorithm runs in time $\tilde{\mathcal{O}}(\min\{n/\varepsilon, n+1/\varepsilon^2\})$ [18, 12, 15].

2 Preliminaries

All logarithms (\log) in this paper are base 2.

Sums and Sumsets For a set S of integers we denote the sum of its elements by $\Sigma(S) := \sum_{s \in S} s$. Given a set Z of n positive integers and target t , the SUBSETSUM problem asks whether there exists $Y \subseteq Z$ with $\Sigma(Y) = t$. We often solve the more general problem of computing the set of all subset sums of Z bounded by t , i.e.,

$$\mathcal{S}(Z; t) := \{\Sigma(Y) \mid Y \subseteq Z\} \cap \{0, \dots, t\}.$$

We represent a set S of integers in $\{0, \dots, m\}$ by its characteristic vector of length $m + 1$. For sets A, B of non-negative integers, we define their *sumset*, in a slightly non-standard way, as the set of all sums of at most one element of A and at most one element of B , i.e.,

$$A \oplus B = \{a + b \mid a \in A \cup \{0\}, b \in B \cup \{0\}\}.$$

For any $t > 0$, we define the *t-capped sumset* as $A \oplus_t B := (A \oplus B) \cap \{0, \dots, t\}$.

Sumset Computation For Boolean vectors x, y of length t we define their convolution as the Boolean vector z of length $2t$ with $z_i = \bigvee_{j=1}^i x_j \wedge y_{i-j}$, where out-of-bounds values are interpreted as false. Observe that the convolution of the characteristic vectors of integer sets A, B equals the characteristic vector of the sumset $A \oplus B$ (if $0 \in A, B$). Thus, sumset computation can be reduced to Boolean convolution.

A simple algorithm for Boolean convolution is to reduce to integer multiplication: Convert the Boolean vector x to a number by interpreting true as 1 and false as 0 and padding with $\log t$ zeroes between any two bits of x . Do the same with y . Then from the product of the constructed numbers we can infer z , since a block of $1 + \log t$ bits is identically 0 if and only if the corresponding bit of z is 0. Hence, if we can multiply two t -bit numbers in time $M(t)$, then Boolean convolution and thus sumset computation can be performed in time $\mathcal{O}(M(t \log t))$. The best known bound for $M(t)$ depends on the machine model. For simplicity, in this paper we work on the RAM model with cell size $w = \Theta(\log t)$, where typical operations on $\log(t)$ -bit numbers can be performed in constant time. In this model it is known that $M(t) = \mathcal{O}(t)$, see, e.g., [10]⁴. This yields the following:

Proposition 2.1. *Given sets A, B of non-negative integers, the t -capped sumset $A \oplus_t B$ can be computed in time $\mathcal{O}(t \log t)$.*

Preprocessing Multisets So far we assumed that the input Z is a *set*. Note that it also makes sense to allow Z to be a *multi-set*, and define a subset $Y \subseteq Z$ to be a multi-set where each $z \in Z$ has multiplicity in Y at most its multiplicity in Z . Thus, any $z \in Z$ with multiplicity m may be used at most m times in any subset sum. E.L. Lawler [18] showed that the general case of multi-sets can be reduced to multi-sets with multiplicities bounded by 2. To this end, for any $z \in Z$ with multiplicity $2k + 1$ we reduce its multiplicity to 1 and add k times the integer $2z$ to Z (i.e. increase the multiplicity of $2z$ by k). Observe that the resulting set generates the same subset sums, since we have $\{i \cdot z \mid 0 \leq i \leq 2k + 1\} = \{i \cdot z + j \cdot 2z \mid 0 \leq i \leq 1 \text{ and } 0 \leq j \leq k\}$. Similarly, if z has multiplicity $2k + 2$ we reduce its multiplicity to 2 and add k times $2z$. Repeating this step for all $z \in Z$ (in increasing order) yields an equivalent SUBSETSUM instance where all multiplicities are bounded by 2, so we obtain the following. For a multi-set Z , by the *size* $|Z|$ we denote the sum of all multiplicities of elements in Z .

Proposition 2.2. *Given a SUBSETSUM instance (Z, t) , where Z is a multi-set of size n , we can in time $\mathcal{O}(n + t)$ compute an equivalent instance (Z', t) where Z' is a multi-set with multiplicities bounded by 2 and $|Z'| \leq \min\{n, 2t\}$.*

Koiliaris and Xu [17] extended this preprocessing as follows. Write the multi-set Z' given by Proposition 2.2 as a union of two *sets* $Z_1, Z_2 \subseteq \{0, \dots, t\}$. Compute $\mathcal{S}(Z_1; t)$ and $\mathcal{S}(Z_2; t)$. Then we obtain $\mathcal{S}(Z'; t)$ as $\mathcal{S}(Z_1; t) \oplus_t \mathcal{S}(Z_2; t)$. This yields a reduction to sets, running in time $\mathcal{O}(n + t \log t)$. A disadvantage is that we have to compute the whole sets $\mathcal{S}(Z_1; t), \mathcal{S}(Z_2; t)$ instead of just solving a decision problem, but this is irrelevant for the algorithms presented in this paper.

By running this reduction as a preprocessing of all our algorithms, throughout the paper we can assume that the input $Z \subseteq \{0, \dots, t\}$ is a set.

3 Near-Linear Time Algorithm

Our goal is to compute, given a set Z of n positive integers and target t , the set $\mathcal{S}(Z; t) := \{\Sigma(Y) \mid Y \subseteq Z\} \cap \{0, \dots, t\}$, since checking whether $t \in \mathcal{S}(Z; t)$ decides the given SUBSETSUM instance. In this section, we design a simple algorithm that solves SUBSETSUM in time $\tilde{\mathcal{O}}(n + t)$, proving Theorem 1.1. As discussed in Section 1.1, our algorithm consists of two parts: We first show how to find sums generated by small subsets using color-coding in Section 3.1, and then use a two-stage approach on layers for finding all subset sums in Section 3.2.

⁴On other machine models $M(t)$ can be larger, e.g., on 2-tape Turing machines or as circuits we have $M(t) \leq t \log t \cdot 2^{\mathcal{O}(\log^* t)}$ [9, 8].

3.1 Color-Coding: An algorithm for small solution size

We describe an algorithm **ColorCoding** (see below) for solving SUBSETSUM if the solution size is small, i.e., an algorithm that finds all sums $\Sigma(Y) \leq t$ generated by sets $Y \subseteq Z$ of size $|Y| \leq k$, for some given (small) k . We *randomly partition* $Z = Z_1 \cup \dots \cup Z_{k^2}$, i.e., we assign any $z \in Z$ to a set Z_i where i is chosen independently and uniformly at random in $\{1, \dots, k^2\}$. We say that this random partition *splits* Y if $|Y \cap Z_i| \leq 1$ holds for all $1 \leq i \leq k^2$. If this happens, then the sumset $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$ contains $\Sigma(Y)$. Indeed, by definition of \oplus in each position i we can choose a number in $Z_i \cup \{0\}$, so for $|Y \cap Z_i| = 1$ we can choose the unique number in the set $Y \cap Z_i$, while for $|Y \cap Z_i| = 0$ we can choose 0, to generate $\Sigma(Y)$ as a sum in $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$. Also note that the sumset $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$ only contains valid sums in $\mathcal{S}(Z; t)$, since no $z \in Z$ may be used twice.

Repeating this procedure sufficiently often with fresh randomness, and taking the union over all computed sumsets $Z_1 \oplus_t \dots \oplus_t Z_{k^2}$, yields a set $S \subseteq \mathcal{S}(Z; t)$ containing any $\Sigma(Y) \leq t$ with $Y \subseteq Z$ and $|Y| \leq k$ with probability at least $1 - \delta$. In fact, $\mathcal{O}(\log 1/\delta)$ repetitions are sufficient, since a random partition splits Y with constant probability, which follows from the birthday paradox, or more precisely the tightness of the birthday paradox bound. We briefly prove this standard claim for completeness. Since $Z \subseteq Z_1 \oplus_t \dots \oplus_t Z_{k^2}$, we can assume $k \geq 2$. For any $Y \subseteq Z$ with $|Y| \leq k$, the probability of the random partition splitting Y is the same as the probability of $|Y|$ balls falling into $|Y|$ different bins, when throwing $|Y|$ balls into k^2 bins. This is equivalent to the second ball falling into a different bin than the first one, the third ball falling into a different bin than the first two, and so on, which has probability

$$\frac{k^2 - 1}{k^2} \cdot \frac{k^2 - 2}{k^2} \cdots \frac{k^2 - (|Y| - 1)}{k^2} \geq \left(\frac{k^2 - (|Y| - 1)}{k^2} \right)^{|Y|} \geq \left(1 - \frac{1}{k} \right)^k \geq \left(\frac{1}{2} \right)^2 = \frac{1}{4}.$$

Hence, $r := \lceil \log_{4/3}(1/\delta) \rceil$ repetitions yield the desired success probability of $1 - (1 - 1/4)^r \geq 1 - \delta$. This finishes the analysis of **ColorCoding** and proves the following lemma. For the running time, note that we perform $\mathcal{O}(\log 1/\delta)$ repetitions of computing k^2 sumsets, each taking time $\mathcal{O}(t \log t)$.

Lemma 3.1. ***ColorCoding**(Z, t, k, δ) computes in time $\mathcal{O}(tk^2 \log t \log(1/\delta))$ a set $S \subseteq \mathcal{S}(Z; t)$ such that for any $Y \subseteq Z$ with $|Y| \leq k$ and $\Sigma(Y) \leq t$ we have $\Sigma(Y) \in S$ with probability $\geq 1 - \delta$.*

Algorithm 1 **ColorCoding**(Z, t, k, δ): Given a set Z of positive integers, target t , size bound $k \geq 1$ and error probability $\delta > 0$, we solve SUBSETSUM with solution size at most k

- 1: **for** $j = 1, \dots, \lceil \log_{4/3}(1/\delta) \rceil$ **do**
 - 2: randomly partition $Z = Z_1 \cup \dots \cup Z_{k^2}$
 - 3: $S_j := Z_1 \oplus_t \dots \oplus_t Z_{k^2}$
 - 4: **return** $\bigcup_j S_j$
-

Standard techniques allow to derandomize this algorithm by iterating over a (deterministic) family of partitions of Z that is guaranteed to contain a partition splitting Y . This comes at the cost of an increased polynomial factor in k , e.g., [20, Lemma 2] gives a factor $\mathcal{O}(k^6 \log k)$. Alternatively, Koiliaris and Xu provide a very different algorithm running in deterministic time $\mathcal{O}(tk^2 \log(tk) \log n)$ [17, Lemma 2.12].

3.2 Layer Splitting

Let (Z, t) be a SUBSETSUM instance and $|Z| = n$. For $\ell \geq 1$, we call (Z, t) an ℓ -*layer* instance if

$$Z \subseteq [t/\ell, 2t/\ell] \quad \text{or} \quad Z \subseteq [0, 2t/\ell] \quad \text{and} \quad \ell \geq n.$$

In both cases we have $Z \subseteq [0, 2t/\ell]$. Moreover, observe that any $Y \subseteq Z$ summing to at most t has size $|Y| \leq \ell$. In the second case, this holds since $|Y| \leq |Z| = n$. Thus, the **ColorCoding** algorithm from the last section with size bound $k = \ell$ solves ℓ -layer instances. In this section, we will show that the running time of **ColorCoding** can be improved for layers, thereby essentially removing the quadratic dependence on ℓ entirely.

However, before discussing how to solve **SUBSETSUM** on layers, we show that any given instance (Z, t) can be split into $\mathcal{O}(\log n)$ layers, see Algorithm 2 below. We simply split set Z at $t/2^i$ for $i = 1, \dots, \lceil \log n \rceil - 1$; this yields $\mathcal{O}(\log n)$ layers $Z_1, \dots, Z_{\lceil \log n \rceil}$. On each layer we then run the algorithm **ColorCodingLayer** presented below, and we combine the resulting sumsets S_i in a straight-forward way. The error probabilities of the calls to **ColorCodingLayer** are chosen sufficiently small so that they sum up to at most δ . Correctness and the running time bound $\mathcal{O}(t \log t \log^3(n/\delta) \log n)$ immediate follow from Lemma 3.2 below. This yields algorithm **FasterSubsetSum** below that proves Theorem 1.1.

Algorithm 2 **FasterSubsetSum** (Z, t, δ) : Returns a set $S \subseteq \mathcal{S}(Z; t)$ containing any $s \in \mathcal{S}(Z; t)$ with probability at least $1 - \delta$, and runs in time $\mathcal{O}(t \log t \log^3(n/\delta) \log n)$

```

1: split  $Z$  into  $Z_i := Z \cap (t/2^i, t/2^{i-1}]$  for  $i = 1, \dots, \lceil \log n \rceil - 1$ , and  $Z_{\lceil \log n \rceil} := Z \cap [0, t/2^{\lceil \log n \rceil - 1}]$ 
2:  $S = \emptyset$ 
3: for  $i = 1, \dots, \lceil \log n \rceil$  do
4:    $S_i := \text{ColorCodingLayer}(Z_i, t, 2^i, \delta/\lceil \log n \rceil)$ 
5:    $S := S \oplus_t S_i$ 
6: return  $S$ 

```

It remains to design a fast algorithm given an ℓ -layer instance (Z, t) and error probability δ . Let $m := \ell/\log(\ell/\delta)$ rounded up to the next power of 2. We randomly partition Z into subsets Z_1, \dots, Z_m . For each Z_j we run **ColorCoding** with size bound $k = 6 \log(\ell/\delta)$, target $12 \log(\ell/\delta)t/\ell$, and error probability δ/ℓ , yielding a set S_j . We combine the sets S_1, \dots, S_m in a natural, binary-tree-like way by computing $S_1 \oplus S_2, S_3 \oplus S_4, \dots, S_{m-1} \oplus S_m$ in the first round, $S_1 \oplus S_2 \oplus S_3 \oplus S_4, \dots, S_{m-3} \oplus S_{m-2} \oplus S_{m-1} \oplus S_m$ in the second round, and so on, until we reach the set $S_1 \oplus \dots \oplus S_m$. Note that in the h -th round of this procedure we combine 2^h sets S_j , initially containing integers bounded by $12 \log(\ell/\delta)t/\ell$. Thus, we may use $\oplus_{2^h \cdot 12 \log(\ell/\delta)t/\ell}$ in the h -th round. This explains the following algorithm.

Algorithm 3 **ColorCodingLayer** (Z, t, ℓ, δ) : See Lemma 3.2 for guarantees

```

1: if  $\ell < \log(\ell/\delta)$  then return ColorCoding $(Z, t, \ell, \delta)$ 
2:  $m := \ell/\log(\ell/\delta)$  rounded up to the next power of 2
3: randomly partition  $Z = Z_1 \cup \dots \cup Z_m$ 
4:  $\gamma := 6 \log(\ell/\delta)$ 
5: for  $j = 1, \dots, m$  do
6:    $S_j := \text{ColorCoding}(Z_j, 2\gamma t/\ell, \gamma, \delta/\ell)$ 
7: for  $h = 1, \dots, \log m$  do ▷ combine the sumsets  $S_j$  in a binary-tree-like way
8:   for  $j = 1, \dots, m/2^h$  do
9:      $S_j := S_{2j} \oplus_{2^h \cdot 2\gamma t/\ell} S_{2j+1}$ 
10: return  $S_1 \cap \{0, \dots, t\}$ 

```

Lemma 3.2. *For an ℓ -layer instance (Z, t) and $\delta \in (0, 1/4]$, the method `ColorCodingLayer` (Z, t, ℓ, δ) computes in time $\mathcal{O}(t \log t \log^3(\ell/\delta))$ a set $S \subseteq \mathcal{S}(Z; t)$ containing any $s \in \mathcal{S}(Z; t)$ with probability at least $1 - \delta$.*

Proof. The case $\ell < \log(\ell/\delta)$ follows from Lemma 3.1. The inclusion $S \subseteq \mathcal{S}(Z; t)$ also follows from Lemma 3.1 and since we only compute sumsets over partitionings, using the fact $(\mathcal{S}(Z_1; t_1) \oplus_{t'} \mathcal{S}(Z_2; t_2)) \cap \{0, \dots, t\} \subseteq \mathcal{S}(Z; t)$ for a partitioning $Z = Z_1 \cup Z_2$ and any $t, t', t_1, t_2 \geq 1$.

For the error probability, fix a subset $Y \subseteq Z$ with $\Sigma(Y) \leq t$, and let $Y_j := Y \cap Z_j$ for $1 \leq j \leq m$. A crucial property is that with sufficiently high probability the size $|Y_j|$ is smaller than $6 \log(\ell/\delta)$, thus allowing us to run `ColorCoding` with size bound $k = 6 \log(\ell/\delta)$.

Claim 3.3. *We have $\Pr[|Y_j| \geq 6 \log(\ell/\delta)] \leq \delta/\ell$.*

Proof. Note that $|Y_j|$ is distributed as the sum of $|Y|$ independent Bernoulli random variables with success probability $1/m$. In particular, $\mu := \mathbb{E}[|Y_j|] = |Y|/m$. A standard Chernoff bound yields that $\Pr[|Y_j| \geq \lambda] \leq 2^{-\lambda}$ for any $\lambda \geq 2e\mu$. Recall that (Z, t) is an ℓ -layer instance, and thus Y has size at most ℓ . This allows us to bound $\mu = |Y|/m \leq \ell/m \leq \log(\ell/\delta)$, by definition of m . The concentration inequality thus holds for $\lambda = 6 \log(\ell/\delta)$, and we obtain $\Pr[|Y_j| \geq 6 \log(\ell/\delta)] \leq \delta/\ell$. \square

By the above claim, we may assume that $|Y_j| \leq 6 \log(\ell/\delta)$ holds for each $1 \leq j \leq m$; this happens with probability at least $1 - m \cdot \delta/\ell$. Since $Z \subseteq [0, 2t/\ell]$, any subset of Z_j of size at most $6 \log(\ell/\delta)$ has sum at most $12 \log(\ell/\delta)t/\ell$. It follows that the call `ColorCoding` $(Z_j, 12 \log(\ell/\delta) \cdot t/\ell, 6 \log(\ell/\delta), \delta/\ell)$ finds $\Sigma(Y_j)$ with probability at least $1 - \delta/\ell$. Assume that this event holds for each $1 \leq j \leq m$; this happens with probability at least $1 - m \cdot \delta/\ell$. Then S_j contains $\Sigma(Y_j)$, and the tree-like sumset computation indeed yields a set containing $\Sigma(Y_1) + \dots + \Sigma(Y_m) = \Sigma(Y)$. The total error probability is $2m\delta/\ell$. Since $\ell \geq 1$ and $\delta \leq 1/4$ we have $\log(\ell/\delta) \geq 2$ and obtain $m \leq \ell/2$. Hence, the total error probability is bounded by δ .

Regarding the running time, observe that by Lemma 3.1 each call to `ColorCoding` takes time $\mathcal{O}(t/\ell \cdot \log^4(\ell/\delta) \log t)$. Since there are $m = \Theta(\ell/\log(\ell/\delta))$ calls to `ColorCoding`, we obtain the claimed time $\mathcal{O}(t \log^3(\ell/\delta) \log t)$. The remaining time for combining the sets S_1, \dots, S_m is

$$\mathcal{O}\left(\sum_{h=1}^{\log m} \frac{m}{2^h} \cdot 2^h \log(\ell/\delta) t/\ell \cdot \log t\right) = \mathcal{O}(t \log t \log m),$$

which is dominated by the total time for calling `ColorCoding`. \square

4 Polynomial Space Algorithm

4.1 Setup

We first give the setup of Lokshtanov and Nederlof [19] and then present a version of their main result using modular arithmetic instead of complex numbers.

Circuits For a set F and binary operators O_1, O_2 on F , a circuit C over $(F; O_1, O_2)$ is a directed acyclic graph $D = (N, A)$ with parallel arcs, such that every node of D is either a constant gate (indegree 0), O_1 gate (indegree 2), or O_2 gate (indegree 2). For an indegree 2 node, its two in-neighbours are its children. The output of a constant gate is the element it is labeled with. The output of an O_i gate is the result of performing O_i on the output of its two children. The output

of C , denoted by $\text{out}(C)$, is the output of a specific gate c marked as the output gate. The size of C is the size of the underlying graph (number of vertices plus number of edges). With some abuse of notation we will denote a gate of C and the output of that gate by the same symbol. A *probabilistic* circuit is a distribution \mathcal{C} over circuits over $(F; O_1, O_2)$.

Vector Operations Fix a ring R with operations $+, \cdot$. For a vector $\mathbf{a} \in R^t$ we write $\mathbf{a}[i]$, $0 \leq i < t$, for its i -th entry. We call \mathbf{a} a *singleton* if it has at most one non-zero entry. We write \boxplus for pointwise addition and \boxtimes for pointwise multiplication, i.e., for vectors $\mathbf{a}, \mathbf{b} \in R^t$ we have $(\mathbf{a} \boxplus \mathbf{b})[i] = \mathbf{a}[i] + \mathbf{b}[i]$ and $(\mathbf{a} \boxtimes \mathbf{b})[i] = \mathbf{a}[i] \cdot \mathbf{b}[i]$. We write \boxtimes for convolution, i.e., we have $(\mathbf{a} \boxtimes \mathbf{b})[i] = \sum_{j=0}^i \mathbf{a}[j] \cdot \mathbf{b}[i-j]$. The *length* of a vector \mathbf{a} is the maximal index of any non-zero entry of \mathbf{a} , i.e., $\text{len}(\mathbf{a}) := \max\{0 \leq i < t \mid \mathbf{a}[i] \neq 0\}$ or 0, if \mathbf{a} is the all-zeroes-vector. We say that the convolution $\mathbf{a} \boxtimes \mathbf{b}$ *overflows* if $\text{len}(\mathbf{a}) + \text{len}(\mathbf{b}) \geq t$. If $\mathbf{a} \boxtimes \mathbf{b}$ does not overflow and \mathbf{a}, \mathbf{b} are the coefficients of polynomials $P(x), Q(x)$, respectively, then $\mathbf{a} \boxtimes \mathbf{b}$ are the coefficients of the polynomial $P(x) \cdot Q(x)$. Throughout the paper we will always work with non-overflowing convolutions. Finally, we also consider the sum of all entries of \mathbf{a} , i.e., $\text{sum}(\mathbf{a}) := \sum_{i=0}^{t-1} \mathbf{a}[i]$.

Lemma 4.1. *Let $\mathbf{a}_1, \dots, \mathbf{a}_k \in R^t$. We have*

- (1) $\text{len}(\mathbf{a}_1 \boxplus \dots \boxplus \mathbf{a}_k) = \max_i \text{len}(\mathbf{a}_i)$,
- (2) $\text{sum}(\mathbf{a}_1 \boxplus \dots \boxplus \mathbf{a}_k) = \sum_i \text{sum}(\mathbf{a}_i)$,
- (3) *if there is no overflow then* $\text{len}(\mathbf{a}_1 \boxtimes \dots \boxtimes \mathbf{a}_k) = \sum_i \text{len}(\mathbf{a}_i)$, *and*
- (4) *if there is no overflow then* $\text{sum}(\mathbf{a}_1 \boxtimes \dots \boxtimes \mathbf{a}_k) = \prod_i \text{sum}(\mathbf{a}_i)$.

Proof. (1), (2), and (3) are immediate. For (4), note that any product $\mathbf{a}_1[j_1] \cdot \dots \cdot \mathbf{a}_k[j_k]$, with $0 \leq j_i \leq \text{len}(\mathbf{a}_i)$, contributes to exactly one entry of $\mathbf{b} := \mathbf{a}_1 \boxtimes \dots \boxtimes \mathbf{a}_k$, and each entry of \mathbf{b} can be written as a sum of such products. \square

Discrete Fourier Transform Fix a prime p and $t \geq 1$. Let $\omega \in \mathbb{Z}_p$ be a primitive t -th root of unity, i.e., $\omega^t = 1$ and $\omega^k \neq 1$ for $1 \leq k < t$. The t -point Discrete Fourier Transform (DFT) over modular arithmetic in \mathbb{Z}_p is the linear function \mathcal{F} mapping a vector $\mathbf{a} \in \mathbb{Z}_p^t$ to $\mathcal{F}(\mathbf{a}) \in \mathbb{Z}_p^t$ with

$$\mathcal{F}(\mathbf{a})[i] = \sum_{j=0}^{t-1} \omega^{ij} \mathbf{a}[j].$$

In other words, if \mathbf{a} is the vector of coefficients of the polynomial $P(x) = \sum_{i=0}^{t-1} \mathbf{a}[i]x^i$ then \mathcal{F} evaluates $P(x)$ at all powers of ω , i.e., $\mathcal{F}(\mathbf{a}) = (P(1), P(\omega), P(\omega^2), \dots, P(\omega^{t-1}))$. The inverse function \mathcal{F}^{-1} is given by

$$\mathcal{F}^{-1}(\mathbf{a})[j] = \frac{1}{t} \sum_{i=0}^{t-1} \omega^{-ij} \mathbf{a}[i].$$

That \mathcal{F}^{-1} is indeed the inverse operation of \mathcal{F} is a consequence of the identity $\sum_{i=0}^{t-1} \omega^{-ij} \omega^{ik} = t \cdot [j = k]$, which follows from ω being a primitive t -th root of unity. Here and in the remainder we use Iverson's bracket notation, i.e., $[B]$ is 1 if B is true, and 0 otherwise. The most important property of DFT for our purposes is the Convolution Theorem (see, e.g., [6]), stating that for any $\mathbf{a}, \mathbf{b} \in \mathbb{Z}_p^t$ such that $\mathbf{a} \boxtimes \mathbf{b}$ does not overflow we have

$$\mathcal{F}(\mathbf{a} \boxtimes \mathbf{b}) = \mathcal{F}(\mathbf{a}) \boxtimes \mathcal{F}(\mathbf{b}).$$

Modular Variant of Lokshtanov-Nederlof Lokshtanov and Nederlof [19] mention without proof that a version of their main result holds for modular arithmetic. For completeness, we prove such a result here.

Theorem 4.2. *Let p be prime, $t \geq 1$, and suppose that \mathbb{Z}_p^t contains a t -th root of unity ω . Let C be a circuit over $(\mathbb{Z}_p^t, \boxplus, \boxtimes)$ with only singleton constants. Suppose that no convolution gate overflows. Then given p, t, ω , and $0 \leq x < t$ we can compute $\text{out}(C)[x]$ in time $\tilde{O}(|C|t \log p)$ and space $\mathcal{O}(|C| \log p)$.*

We discuss how to find an appropriate root of unity ω later in Lemma 4.5. We remark that Lokshtanov and Nederlof have the depth of C as an additional factor in their time and space bounds, which seems to be due to their choice of the complex Fourier transform and does not appear in the modular version (this is a lower-order improvement).

Proof. The proof by Lokshtanov and Nederlof works almost verbatim. From the circuit C we construct another circuit C' over $(\mathbb{Z}_p^t, \boxplus, \boxdot)$ with the same directed graph as C , but with different gates. Each constant gate $\mathbf{a} \in C$ is replaced with the constant gate $\mathbf{a}' = \mathcal{F}(\mathbf{a})$. Each convolution gate \boxtimes is replaced with a pointwise multiplication gate \boxdot . Then an easy inductive argument shows that for any gate $\mathbf{a} \in C$ and its corresponding gate $\mathbf{a}' \in C'$ we have $\mathbf{a}' = \mathcal{F}(\mathbf{a})$. Indeed, if \mathbf{a} is a constant gate then the claim holds by definition. For the inductive step, consider let \mathbf{b}, \mathbf{c} be the children of \mathbf{a} . By the inductive hypothesis we have $\mathbf{b}' = \mathcal{F}(\mathbf{b})$ and $\mathbf{c}' = \mathcal{F}(\mathbf{c})$. If \mathbf{a} is an addition gate, then we conclude $\mathbf{a}' = \mathbf{b}' \boxplus \mathbf{c}' = \mathcal{F}(\mathbf{b}) \boxplus \mathcal{F}(\mathbf{c}) = \mathcal{F}(\mathbf{b} \boxplus \mathbf{c}) = \mathcal{F}(\mathbf{a})$. If \mathbf{a} is a convolution gate, then the Convolution Theorem implies $\mathbf{a}' = \mathbf{b}' \boxdot \mathbf{c}' = \mathcal{F}(\mathbf{b}) \boxdot \mathcal{F}(\mathbf{c}) = \mathcal{F}(\mathbf{b} \boxtimes \mathbf{c}) = \mathcal{F}(\mathbf{a})$.

For ease of notation we write $\mathbf{f} := \text{out}(C)$. We now use C' to compute $\mathbf{f}[x]$. Plugging the definition of \mathcal{F}^{-1} into the identity $\mathbf{f} = \mathcal{F}^{-1}(\mathcal{F}(\mathbf{f}))$ yields

$$\mathbf{f}[x] = \frac{1}{t} \sum_{i=0}^{t-1} \omega^{-ix} (\mathcal{F}(\mathbf{f}))[i].$$

Note that ω^s , for $s \leq \text{poly}(t)$, can be computed with $\mathcal{O}(\log t)$ arithmetic operations in \mathbb{Z}_p by repeated squaring. Thus, in order to compute $\mathbf{f}[x]$ it suffices to compute $(\mathcal{F}(\mathbf{f}))[0], \dots, (\mathcal{F}(\mathbf{f}))[t-1]$.

For computing $(\mathcal{F}(\mathbf{f}))[j]$, we convert circuit C' to a circuit C'_j over $(\mathbb{Z}_p, +, \cdot)$ with the same directed graph as C' . Each constant gate $\mathbf{a}' \in C'$ is replaced by $\mathbf{a}'[j]$. Each pointwise addition gate \boxplus is replaced with $+$ and each pointwise multiplication gate \boxdot is replaced with \cdot . It follows immediately that C'_j computes $(\mathcal{F}(\mathbf{f}))[j]$.

Consider a (singleton) constant gate $\mathbf{a} \in C$ with $\mathbf{a}[k] = v$ and $\mathbf{a}[\ell] = 0$ for all $\ell \neq k$. By definition of \mathcal{F} we have $(\mathcal{F}(\mathbf{a}))[j] = \omega^{jk}v$. Thus, any constant gate in C'_j can be computed with $\mathcal{O}(\log t)$ arithmetic operations in \mathbb{Z}_p . Computing the remaining gates of C'_j takes $|C'| = |C|$ arithmetic operations in \mathbb{Z}_p . As arithmetic operations in \mathbb{Z}_p can be performed in time $\tilde{O}(\log p)$, we can compute any value $(\mathcal{F}(\mathbf{a}))[j]$ in time $\tilde{O}(|C| \log p \log t)$ and space $\mathcal{O}(|C| \log p)$. Summing over all $j = 0, \dots, t-1$ yields time $\tilde{O}(|C|t \log p \log t) = \tilde{O}(|C|t \log p)$ and space $\mathcal{O}(|C| \log p)$. \square

4.2 FasterSubsetSum as a Circuit

In the remainder of the paper, set the error probability $\delta := 1/n$. We now describe how to convert the algorithm **FasterSubsetSum** into a probabilistic circuit over $(\mathbb{N}^{\tilde{O}(t)}; \boxplus, \boxtimes)$. This yields:

Lemma 4.3. *Given a SUBSETSUM instance (Z, t) with $|Z| = n$, we can construct a probabilistic circuit C over $(\mathbb{N}^{t'}; \boxplus, \boxtimes)$, where $t' = \mathcal{O}(t \log^3 n)$ is a power of two, satisfying*

(1) for any $0 \leq s \leq t$, $s \notin \mathcal{S}(Z; t)$ we have $\Pr_{C \in \mathcal{C}}[\text{out}(C)[s] = 0] = 1$, and

(2) for any $0 \leq s \leq t$, $s \in \mathcal{S}(Z; t)$ we have $\Pr_{C \in \mathcal{C}}[\text{out}(C)[s] = 0] \leq 1/n$.

Any circuit $C \in \mathcal{C}$ has size $|C| = \tilde{O}(n)$, uses only singleton constants, has no overflowing convolution gates, and any gate $\mathbf{a} \in C$ satisfies $\text{sum}(\mathbf{a}) \leq 2^{\tilde{O}(n)}$. Sampling a circuit C from \mathcal{C} can be performed in time $\tilde{O}(n)$.

Proof. We carefully inspect the algorithm **FasterSubsetSum**, transforming any \cup -operation into a \boxplus -gate and any \oplus_t -operation into a \boxtimes -gate. The details are as follows.

Set $t' := 300 \log^3(2n^2)t$, rounded up to a power of two. For any vector $\mathbf{a} \in \mathbb{N}^{t'}$ we say that \mathbf{a} represents the set $\{0 \leq i < t' \mid \mathbf{a}[i] > 0\}$. Note that if \mathbf{a}, \mathbf{b} represent sets A, B then $\mathbf{a} \boxplus \mathbf{b}$ represents $A \cup B$. Moreover, if $\mathbf{a} \boxtimes \mathbf{b}$ does not overflow then $\mathbf{a} \boxtimes \mathbf{b}$ represents $A \oplus_{t'} B$, which in this case equals $A \oplus B$. For any $0 \leq i < t'$ we denote by \mathbf{e}_i the singleton vector containing a 1 at position i and a 0 at all other entries. Note that \mathbf{e}_i represents $\{i\}$.

We now use this representation to convert algorithm **FasterSubsetSum** to a probabilistic circuit over $(\mathbb{N}^{t'}; \boxplus, \boxtimes)$. Fix the randomness of **FasterSubsetSum**, i.e., fix all random partitions performed in **ColorCoding** and **ColorCodingLayer**. All remaining basic operations in **FasterSubsetSum**, **ColorCoding**, and **ColorCodingLayer** are (1) forming a subset $Z' \subseteq Z$, (2) taking the union of two sets, and (3) computing the sumset of two sets. Hence, we can view **FasterSubsetSum** as a circuit where each node computes a set. Each leaf is directly assigned a subset $Z' \subseteq Z$. Each inner node either computes the union \cup or the sumset $\oplus_{t'}$, for some t'' , of its children.

Now we replace each leaf, corresponding to $Z' \subseteq Z$, by the circuit computing $\boxplus_{z \in Z'} \mathbf{e}_z$, which is a vector representing Z' . Note that this circuit has size $|Z'|$ and uses only singleton constants. Moreover, we replace each \cup -gate by \boxplus and each $\oplus_{t''}$ -gate by \boxtimes . This yields a circuit C over $(\mathbb{N}^{t'}; \boxplus, \boxtimes)$. Over the randomness of **FasterSubsetSum** used for the random partitionings, we obtain a probabilistic circuit \mathcal{C} .

We show that t' is chosen sufficiently large so that circuit C has no overflowing convolution gates. To this end, we use Lemma 4.1 to bound the length $\text{len}(\mathbf{a}) = \max\{0 \leq i < t' \mid \mathbf{a}[i] \neq 0\}$ of any gate \mathbf{a} . For the output of **ColorCoding** (Z, t, k, δ) , as well as any of its inner gates, we can bound the length by $k^2 \cdot \max Z$, since we compute a union (which does not increase the length) over a k^2 -wise sumset computation (multiplying the length by k^2) over subsets of Z (having length at most $\max Z$). For **ColorCodingLayer** (Z, t, ℓ, δ) we can bound the length by $\ell \cdot 36 \log^2(\ell/\delta) \max Z$, since we compute an m -wise sumset, for some $m \leq \ell$, over the output of **ColorCoding** with $k = 6 \log(\ell/\delta)$. Finally, for **FasterSubsetSum** (Z, t, δ) we perform a sumset computation over layers $i = 1, \dots, \lceil \log n \rceil$. For the i -th layer we run **ColorCodingLayer** on $\ell = 2^i$ and Z_i with $\max Z_i \leq t/2^{i-1}$, which by the above bound yields length at most $2^i \cdot 36 \log^2(2^i/\delta) t/2^{i-1} \leq 72 \log^2(2n/\delta) t$. Summing over all $i = 1, \dots, \lceil \log n \rceil$ yields the desired bound on the length of $\lceil \log n \rceil \cdot 72 \log^2(2n/\delta) t \leq 144 \log^3(2n/\delta) t < t'/2$, since we set $\delta = 1/n$. This bound holds at all gates of any circuit $C \in \mathcal{C}$. In particular, this shows that no convolution gate overflows, since t' is chosen sufficiently large.

With this property it is immediate that C computes a vector $\text{out}(C)$ representing a set S' such that for $S := S' \cap \{0, \dots, t\}$ we have **FasterSubsetSum** $(Z, t, 1/n) \subseteq S \subseteq \mathcal{S}(Z; t)$. Indeed, by the same reasoning as for **FasterSubsetSum** we argue that the resulting circuit only performs set partitionings and sumset computations and thus only computes valid subset sums, proving $S' \subseteq \mathcal{S}(Z; t')$ and thus $S \subseteq \mathcal{S}(Z; t)$. For the other direction, note that we drop the index from a $\oplus_{t''}$ -gate when converting it to a \boxtimes -gate. Thus, the circuit C does not compute the same set S as **FasterSubsetSum** $(Z, t, 1/n)$, but S may be a proper superset. However, this change cannot remove elements, so we obtain the inclusion **FasterSubsetSum** $(Z, t, 1/n) \subseteq S$.

From this, correctness follows immediately. Indeed, for $s \in \{0, \dots, t\}$, $s \notin \mathcal{S}(Z; t)$ we obtain $s \notin S$ and thus $\text{out}(C)[s] = 0$. Also, for $s \in \mathcal{S}(Z; t)$ we have $s \in \text{FasterSubsetSum}(Z, t, 1/n)$ with probability at least $1 - 1/n$ and thus $s \in S$ with probability at least $1 - 1/n$.

Since we almost always simply partition the current set Z , except for **ColorCoding** where we take a union over $\mathcal{O}(\log(n))$ partitionings, each item in Z appears in $\mathcal{O}(\log(n))$ leaves of the circuit. It is also easy to see that the depth of the circuit is $\mathcal{O}(\log(n))$. This allows us to bound $|C| \leq \tilde{\mathcal{O}}(n)$, and it is easy to see that $C \in \mathcal{C}$ can be sampled in the same time bound.

Finally, we bound the sum of entries $\text{sum}(\mathbf{a}) = \sum_{i=0}^{t'-1} \mathbf{a}[i]$ for any gate $\mathbf{a} \in C$ with $C \in \mathcal{C}$, using Lemma 4.1.(4). For a subset $Z' \subset Z$ we implement a circuit representing Z' with sum of entries equal to $|Z'| \leq n$. For **ColorCoding** (Z, t, k, δ) we have sum of entries $\mathcal{O}(\log(1/\delta)n^{k^2})$, since we take the union over $\mathcal{O}(\log(1/\delta))$ rounds (multiplying the sum of entries by $\mathcal{O}(\log(1/\delta))$) over k^2 -wise sumset computations (raising the sum of entries to the k^2 -th power) over subsets of Z , which have sum of entries at most n . For **ColorCodingLayer** (Z, t, ℓ, δ) we can bound the sum of entries by $n^{\mathcal{O}(\log^2(\ell/\delta) \cdot \ell)}$, since we compute an m -wise sumset, for some $m \leq \ell$, over the output of **ColorCoding** with $k = \mathcal{O}(\log(\ell/\delta))$. Finally, for **FasterSubsetSum** $(Z, t, 1/n)$ the sum of entries is bounded by $n^{\mathcal{O}(\log n \cdot \log^2(n) \cdot n)} = 2^{\mathcal{O}(\log^4(n) \cdot n)} = 2^{\tilde{\mathcal{O}}(n)}$, since we perform a sumset computation over $\mathcal{O}(\log n)$ layers (which raises the sum of entries to the $\mathcal{O}(\log n)$ -th power) over the result of **ColorCodingLayer** on $\ell \leq 2n$. \square

Note that the bound $\text{sum}(\mathbf{a}) \leq 2^{\tilde{\mathcal{O}}(n)}$ also bounds the maximal entry of any gate \mathbf{a} in any circuit $C \in \mathcal{C}$. Thus, we may replace \mathbb{N} by \mathbb{Z}_p for some $p \leq 2^{\tilde{\mathcal{O}}(n)}$ (with sufficiently many hidden logarithmic factors) and still obtain the same result. This yields a probabilistic circuit C' over $(\mathbb{Z}_p^{t'}, \boxplus, \boxtimes)$. Running Theorem 4.2 on C' would yield a randomized SUBSETSUM algorithm with time $\tilde{\mathcal{O}}(nt \log p) = \tilde{\mathcal{O}}(n^2 t)$ and space $\tilde{\mathcal{O}}(n \log p) = \tilde{\mathcal{O}}(n^2)$, if we were able to efficiently compute the required root of unity (which is a non-trivial problem).

In the next section, we further reduce time and space and we provide an algorithm for computing an appropriate root of unity.

4.3 Reducing the Domain Size

In the last section, we designed a probabilistic circuit \mathcal{C} over $(\mathbb{N}^{t'}; \boxplus, \boxtimes)$ and then replaced \mathbb{N} by \mathbb{Z}_p for sufficiently large $p = 2^{\tilde{\mathcal{O}}(n)}$. Now we will pick p as a random prime, chosen uniformly from a set P of size $\tilde{\Omega}(n^2)$. This randomness is independent from the randomness of picking a circuit $C \in \mathcal{C}$. Since any number $k \leq 2^{\tilde{\mathcal{O}}(n)}$ has at most $\tilde{\mathcal{O}}(n)$ prime factors, for any $C \in \mathcal{C}$ the probability that p divides $\text{out}(C)$ is at most $\tilde{\mathcal{O}}(1/n)$. Interpreting a circuit $C \in \mathcal{C}$ as a circuit over $(\mathbb{Z}_p^{t'}; \boxplus, \boxtimes)$, i.e., performing all arithmetic operations of C modulo p , yields a circuit C_p over $(\mathbb{Z}_p^{t'}; \boxplus, \boxtimes)$ that computes the value $\text{out}(C) \bmod p$. Taking into account the randomness of $C \in \mathcal{C}$, we obtain a probabilistic circuit \mathcal{C}_p over $(\mathbb{Z}_p^{t'}; \boxplus, \boxtimes)$ with the following properties (see Lemma 4.3):

- (1) for any $0 \leq s \leq t$, $s \notin \mathcal{S}(Z; t)$ we have $\Pr_{p \in P}[\Pr_{C_p \in \mathcal{C}_p}[\text{out}(C_p)[s] = 0]] = 1$, and
- (2) for any $0 \leq s \leq t$, $s \in \mathcal{S}(Z; t)$ we have $\Pr_{p \in P}[\Pr_{C_p \in \mathcal{C}_p}[\text{out}(C_p)[s] = 0]] \leq \tilde{\mathcal{O}}(1/n)$.

Note that Theorem 4.2 and Lemma 4.3 now yield total time $\tilde{\mathcal{O}}(|C|t \log(\max P)) = \tilde{\mathcal{O}}(nt \log(\max P))$ and space $\tilde{\mathcal{O}}(|C| \log(\max P)) = \tilde{\mathcal{O}}(n \log(\max P))$ for solving SUBSETSUM.

It remains to choose an appropriate set of primes P . From the above discussion we have the following requirements, where we need (1) for running DFT, see Theorem 4.2, (2) to avoid $\tilde{\mathcal{O}}(n)$ prime factors, and (3) since $\max P$ appears in the running time.

Lemma 4.4. *For any power of two τ , and $n \leq \tau$, there is a set P of primes such that*

- (1) \mathbb{Z}_p contains a τ -th root of unity ω_p for any $p \in P$,
- (2) $|P| = \tilde{\Omega}(n^2)$, and
- (3) assuming ERH we have $\max P = \tau^{\mathcal{O}(1)}$, while unconditionally we have $\max P \leq \mathcal{O}(\exp(\tau^\varepsilon))$ for any $\varepsilon > 0$ (with the $\text{polylog}(n)$ in $|P| = \tilde{\Omega}(n^2)$ depending on ε).

We can pick a random $p \in P$ and compute ω_p in randomized time $(\log(\max P))^{\mathcal{O}(1)}$, with error probability $1/\text{poly}(\tau)$.

Recall from Theorem 4.2 that the time and space bounds incur a factor $\mathcal{O}(\log p)$. Assuming ERH we have $\mathcal{O}(\log p) = \mathcal{O}(\log \tau) = \mathcal{O}(\log t)$, which yields the desired time $\tilde{\mathcal{O}}(nt)$ and space $\tilde{\mathcal{O}}(n \log t)$, proving the first statement of Theorem 1.3. The unconditional statement follows similarly.

Proof. We first simplify requirement (1). For any prime p , by Fermat's little theorem any $a \in \mathbb{Z}_p$, $a \neq 0$, is a $(p-1)$ -th root of unity. The additional property $a^k \neq 1$ for all $1 \leq k < p-1$ of a *primitive* $(p-1)$ -th root of unity is satisfied for any generator g of the multiplicative group $\mathbb{Z}_p^* = \mathbb{Z}_p \setminus \{0\}$. Finally, if τ divides $p-1$ and g is a primitive $(p-1)$ -th root of unity then $g^{(p-1)/\tau}$ is a primitive τ -th root of unity. Note that if τ divides $p-1$ then p is in the arithmetic progression $\{1 + k \cdot \tau \mid k \in \mathbb{N}\}$, which we write as $1 + \tau \cdot \mathbb{N}$. Hence, we may replace requirement (1) by

$$(1') \quad P \subseteq 1 + \tau \cdot \mathbb{N}.$$

The above argument shows existence of a primitive τ -th root of unity, but it does not yield an efficient algorithm for determining one, since there is no efficient algorithm known for finding a generator of \mathbb{Z}_p^* (the known algorithms have to compute a prime factorization of $p-1$). We circumvent this problem by using that τ is a power of two. In this situation, we can compute a τ -th root of unity as follows.

Lemma 4.5. *Let p be prime and τ be a power of two, where τ divides $p-1$. Pick $a \in \mathbb{Z}_p^*$ uniformly at random and set $\omega := a^{(p-1)/\tau}$. If $\omega^\tau \neq 1$ or $\omega^{\tau/2} = 1$ then restart this procedure. This method finds a primitive τ -th root of unity $\omega \in \mathbb{Z}_p$ in expected time $\mathcal{O}(\text{polylog}(p))$.*

Note that by halting after $\mathcal{O}(\text{polylog}(p))$ iterations we may obtain an algorithm with *worst-case* running time $\mathcal{O}(\text{polylog}(p))$ and error probability $1/\text{poly}(p) \leq 1/\text{poly}(t)$.

Proof. First observe that if the method finishes then ω is a primitive τ -th root of unity. Indeed, ω is a τ -th root of unity since $\omega^\tau = 1$. For showing that ω is primitive, consider the minimal $\ell \geq 1$ with $\omega^\ell = 1$. From $\omega^\tau = 1$ it follows that ℓ divides τ , and thus ℓ is a power of two. In particular, if $\ell < \tau$ then $\omega^{\tau/2} = \omega^{\ell \cdot \tau/(2\ell)} = 1^{\tau/(2\ell)} = 1$, contradicting the break condition $\omega^{\tau/2} \neq 1$. Hence, $\ell = \tau$, and for any $1 \leq k < \tau$ we have $\omega^k \neq 1$, showing that ω is primitive.

Since arithmetic operations in \mathbb{Z}_p can be performed in time $\mathcal{O}(\text{polylog}(p))$, it remains to bound the success probability of one round of the method by $\Omega(1/\text{polylog}(p))$. To this end, it suffices to show that a random $a \in \mathbb{Z}_p^*$ is likely to be a generator of \mathbb{Z}_p^* , since then a is a primitive $(p-1)$ -th root of unity and $\omega = a^{(p-1)/\tau}$ is a primitive τ -th root of unity. For bounding the probability of a being a generator, we combine the well-known fact that the number of generators is $\phi(p-1)$, where ϕ is Euler's totient function, and the asymptotic lower bound $\phi(n) = \Omega(n/\log \log n)$. \square

It remains to find a set $P \subseteq 1 + \tau \cdot \mathbb{N}$ of $\tilde{\Omega}(n^2)$ primes. Dirichlet's theorem shows that the arithmetic progression $1 + \tau \cdot \mathbb{N}$ contains infinitely many primes. However, in order to bound $\max P$ we need a quantitative version of Dirichlet's theorem. In particular, let $\pi(x, \tau)$ be the number of primes in $(1 + \tau \cdot \mathbb{N}) \cap [0, x]$. We need an upper bound on the minimal x such that $\pi(x, \tau) = \tilde{\Omega}(n^2)$.

The best unconditional bound is given by (a well-known corollary of) the Siegel-Walfisz theorem [24]: For any constant $C > 0$ there is a constant $C' > 0$ such that if $\tau \leq (\log x)^C$ then

$$\pi(x, \tau) = \text{Li}(x)/\phi(\tau) \pm \mathcal{O}(x \exp(-C' \sqrt{\log x})),$$

where Li denotes the offset logarithmic integral. Using $\text{Li}(x) = \Theta(x/\log x)$ and $\phi(\tau) \leq \tau \leq (\log x)^C$, we obtain $\text{Li}(x)/\phi(\tau) = \tilde{\Omega}(x)$, which dominates the error term $\mathcal{O}(x \exp(-C' \sqrt{\log x}))$ for any sufficiently large x . Hence, we obtain $\pi(x, \tau) = \tilde{\Omega}(x)$ for $\tau \leq (\log x)^C$. Setting $C = 1/\varepsilon$ and rearranging, we have $\pi(x, \tau) = \tilde{\Omega}(x)$ for $x \geq \exp(\tau^\varepsilon)$.

Hence, for the threshold $x := n^2 + \exp(\tau^\varepsilon) = \mathcal{O}(\exp(\tau^\varepsilon))$ we have $\pi(x, \tau) = \tilde{\Omega}(n^2)$. We now set P as the set of primes in $R := (1 + \tau \cdot \mathbb{N}) \cap [0, x]$. Clearly, we have $|P| = \tilde{\Omega}(n^2)$. Moreover, note that the set R contains $(x - 1)/\tau \leq x$ numbers, $\tilde{\Omega}(x)$ of which are prime. Thus, a random number in R is prime with probability $\Omega(1/\text{polylog}(x)) = \tau^{-\mathcal{O}(\varepsilon)}$. Now to pick a random $p \in P$, we repeatedly pick a random number $r \in R$ until it is prime. Recall that checking primality can be done in polynomial time $\mathcal{O}(\text{polylog}(x)) = \tau^{\mathcal{O}(\varepsilon)}$ [2]. The expected number of repetitions until we find a prime is also $\tau^{\mathcal{O}(\varepsilon)}$. By halting after essentially the same number of repetitions we can also obtain a *worst-case* time bound with sufficiently small error probability, say $1/\tau$. The total time for finding a random prime $p \in P$ is thus $\tau^{\mathcal{O}(\varepsilon)}$. This procedure has to be performed only once, so this time is negligible compared to the remaining running time $\tilde{\mathcal{O}}(nt)$. Note that the number of log-factors in the bound $|P| = \tilde{\Omega}(n^2)$ depends on ε , and since the constants in the Siegel-Walfisz theorem are ineffective, we cannot give an explicit bound on the number of log-factors in terms of ε .

Assuming the Extended Riemann Hypothesis, very good asymptotic bounds on $\pi(x, \tau)$ are known. Using a bound by Titchmarsh, in [8, Proposition 4.3] it is shown, along the lines of the last two paragraphs, that already for threshold $x := \tilde{\mathcal{O}}(\tau(\tau + n^2))$ (with sufficiently many hidden logarithmic factors) the set R contains $\tilde{\Omega}(n^2)$ primes, and sampling random elements of R yields a prime in expected time $\mathcal{O}(\text{polylog}(x)) = \mathcal{O}(\text{polylog}(\tau))$. This finishes the proof. \square

5 Unbounded Subset Sum

We briefly discuss the unbounded variant of SUBSETSUM, where each input number can be chosen arbitrarily often (whereas in the bounded variant that we studied so far each input number can be chosen at most once). Specifically, given a set Z of n positive integers and target t , the task is to determine whether any sequence over Z sums to exactly t . Note that here it makes no sense to define Z as a multi-set and thus $n \leq t$. We present a simple $\tilde{\mathcal{O}}(t)$ algorithm.

First consider the problem whether any sequence over Z of length at most k sums to t . Clearly, this problem can be solved in time $\tilde{\mathcal{O}}(kt)$, by computing the k -fold sumset $Z \oplus_t Z \oplus_t \dots \oplus_t Z$, and checking whether it contains t . Note that for the usual bounded variant of SUBSETSUM this property breaks down, since already the sumset $Z \oplus_t Z$ is not necessarily contained in $\mathcal{S}(Z; t)$, as it contains sums of the form $z + z$ for $z \in Z$, so we have to resort to sumsets $Z_1 \oplus_t \dots \oplus_t Z_k$ over partitionings $Z = Z_1 \cup \dots \cup Z_k$. This is a reason why (bounded) SUBSETSUM is much harder than the unbounded version UNBOUNDEDSUBSETSUM.

Algorithm 4 solves the UNBOUNDEDSUBSETSUM problem for each $t' \leq t$ at once, i.e., it computes $\mathcal{S}^{unb}(Z; t) := \{0, \dots, t\} \cap \{a_1 + \dots + a_k \mid k \in \mathbb{N}_{\geq 0}, a_1, \dots, a_k \in Z\}$. We use the fact that the classic dynamic programming algorithm computes $\mathcal{S}^{unb}(Z; t)$ in time $\mathcal{O}(nt)$.

Algorithm 4 UNBOUNDEDSUBSETSUM in time $\tilde{\mathcal{O}}(t)$. Given n positive integers Z and target t , the algorithm computes for each $t' \leq t$ whether a sequence over Z sums to t' .

```

1: compute  $S_0 := \mathcal{S}^{unb}(Z; t/n)$  using the classic dynamic program in time  $\mathcal{O}(n \cdot t/n) = \mathcal{O}(t)$ 
2: for  $i = 1, \dots, \lceil \log n \rceil$  do
3:    $t_i := 2^i t/n$ 
4:    $S_i := S_{i-1} \oplus_{t_i} S_{i-1} \oplus_{t_i} Z$ 
5: return  $S_{\lceil \log n \rceil} \cap \{0, \dots, t\}$ 

```

Since the i -th iteration takes time $\tilde{\mathcal{O}}(t_i) = \tilde{\mathcal{O}}(2^i t/n)$, the total running time of Algorithm 4 is $\tilde{\mathcal{O}}(t)$, more precisely $\mathcal{O}(t \log t)$. For correctness, we argue inductively that $S_i = \mathcal{S}^{unb}(Z; 2^i t/n)$, and thus $S_{\lceil \log n \rceil} \cap \{0, \dots, t\} = \mathcal{S}^{unb}(Z; t)$. Consider any sequence A over Z summing to an integer in $[0, 2^i t/n]$. If $\Sigma(A) \leq 2^{i-1} t/n$, then

$$\Sigma(A) \in \mathcal{S}^{unb}(Z; 2^{i-1} t/n) = S_{i-1} \subseteq S_{i-1} \oplus_{t_i} S_{i-1} \oplus_{t_i} Z = S_i.$$

Thus, assume $\Sigma(A) > 2^{i-1} t/n$. We split $A = (a_1, \dots, a_k)$ at the smallest index j with $a_1 + \dots + a_j > 2^{i-1} t/n$ into $A_1 = (a_1, \dots, a_{j-1})$, a_j , and $A_2 = (a_{j+1}, \dots, a_k)$. Observing that $a_j \in Z$ and $\Sigma(A_1), \Sigma(A_2) \leq 2^{i-1} t/n$ and hence $\Sigma(A_1), \Sigma(A_2) \in S_{i-1}$, we obtain $\Sigma(A) \in S_{i-1} \oplus_{t_i} S_{i-1} \oplus_{t_i} Z = S_i$. Hence, $\mathcal{S}^{unb}(Z; 2^i t/n) \subseteq S_i$. Correctness follows, since clearly the converse $S_i \subseteq \mathcal{S}^{unb}(Z; 2^i t/n)$ holds as well.

6 Conclusion

A textbook algorithm solves SUBSETSUM in pseudopolynomial time $\mathcal{O}(nt)$ on inputs consisting of n positive numbers with target t . As our main result we present an improved algorithm running in time $\tilde{\mathcal{O}}(n + t)$. This improves upon a classic algorithm and is likely to be near-optimal, since it matches conditional lower bounds of $t^{1-\varepsilon}$ for any $\varepsilon > 0$ from SETCOVER and combinatorial k -CLIQUE. Our algorithm heavily uses randomization and has one-sided error. The main tricks are to find *small* solutions using color-coding, and to split *large* solutions by a two-stage color-coding-like process.

We also improve the best known polynomial space algorithm for SUBSETSUM to time $\tilde{\mathcal{O}}(nt)$ and space $\tilde{\mathcal{O}}(n)$ assuming the Extended Riemann Hypothesis, and to time $\tilde{\mathcal{O}}(n t^{1+\varepsilon})$ and space $\tilde{\mathcal{O}}(n t^\varepsilon)$ for any $\varepsilon > 0$ unconditionally. This improvement is achieved by combining our $\tilde{\mathcal{O}}(n + t)$ algorithm with the previously best polynomial space algorithm by Lokshtanov and Nederlof [19], and by working modulo a random prime in an arithmetic progression, which explains the connection to ERH. We leave it as an open problem to obtain an algorithm with time $\tilde{\mathcal{O}}(n + t)$ and space $\tilde{\mathcal{O}}(n)$.

Our techniques are very different from the recent $\tilde{\mathcal{O}}(\sqrt{n}t)$ algorithm by Koiliaris and Xu [17], which is the fastest known deterministic algorithm.

Acknowledgements The author wants to thank Marvin Künnemann and Jesper Nederlof for providing useful comments on a draft of this paper.

References

- [1] A. Abboud, K. Lewi, and R. Williams. Losing weight by gaining edges. In *ESA'14*, pages 1–12, 2014.

- [2] M. Agrawal, N. Kayal, and N. Saxena. PRIMES is in P. *Annals of Mathematics*, pages 781–793, 2004.
- [3] N. Alon, R. Yuster, and U. Zwick. Color-coding. *Journal of the ACM*, 42(4):844–856, 1995.
- [4] P. Austrin, P. Kaski, M. Koivisto, and J. Nederlof. Subset sum in the absence of concentration. In *STACS’15*, pages 48–61, 2015.
- [5] R.E. Bellman. *Dynamic programming*. Princeton University Press, 1957.
- [6] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.
- [7] M. Cygan, H. Dell, D. Lokshtanov, D. Marx, J. Nederlof, Y. Okamoto, R. Paturi, S. Saurabh, and M. Wahlström. On problems as hard as CNF-SAT. *ACM Transactions on Algorithms*, 12(3):41, 2016.
- [8] A. De, P. P. Kurur, C. Saha, and R. Saptharishi. Fast integer multiplication using modular arithmetic. *SIAM Journal on Computing*, 42(2):685–699, 2013.
- [9] M. Fürer. Faster integer multiplication. *SIAM Journal on Computing*, 39(3):979–1005, 2009.
- [10] M. Fürer. How fast can we multiply large integers on an actual computer? In *LATIN’14*, pages 660–670, 2014.
- [11] Z. Galil and O. Margalit. An almost linear-time algorithm for the dense subset-sum problem. *SIAM Journal on Computing*, 20(6):1157–1189, 1991.
- [12] G. V. Gens and E. V. Levner. Fast approximation algorithms for knapsack type problems. In *Optimization Techniques*, pages 185–194. Springer, 1980.
- [13] E. Horowitz and S. Sahni. Computing partitions with applications to the knapsack problem. *Journal of the ACM*, 21(2):277–292, 1974.
- [14] T. Husfeldt, R. Paturi, G.B. Sorkin, and R. Williams. Exponential Algorithms: Algorithms and Complexity Beyond Polynomial Time (Dagstuhl Seminar 13331). *Dagstuhl Reports*, 3(8):40–72, 2013.
- [15] H. Kellerer, R. Mansini, U. Pferschy, and M. G. Speranza. An efficient fully polynomial approximation scheme for the subset-sum problem. *Journal of Computer and System Sciences*, 66(2):349–370, 2003.
- [16] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack problems*. Springer, 2004.
- [17] K. Koiliaris and C. Xu. A faster pseudopolynomial time algorithm for subset sum. In *SODA’17*, 2017. To appear. Preprint available at arxiv.org/abs/1507.02318.
- [18] E. L. Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 4(4):339–356, 1979.
- [19] D. Lokshtanov and J. Nederlof. Saving space by algebraization. In *STOC’10*, pages 321–330, 2010.
- [20] M. Naor, L. J. Schulman, and A. Srinivasan. Splitters and near-optimal derandomization. In *FOCS’95*, pages 182–191, 1995.

- [21] D. Pisinger. Linear time algorithms for knapsack problems with bounded weights. *Journal of Algorithms*, 33(1):1–14, 1999.
- [22] D. Pisinger. Dynamic programming on the word RAM. *Algorithmica*, 35(2):128–145, 2003.
- [23] J. P. Schmidt and A. Siegel. The spatial complexity of oblivious k-probe hash functions. *SIAM Journal on Computing*, 19(5):775–786, 1990.
- [24] A. Walfisz. Zur additiven Zahlentheorie. II. *Mathematische Zeitschrift*, 40(1):592–607, 1936.